

A Brief Guide to Advection Operators in FV3
Lucas Harris, Shian-Jiann Lin, and Xi Chen, GFDL
17 August 2018

FV3 is unique among atmospheric dynamical cores in that it is formulated similarly to CFD solvers which treat most processes as the fluxes of some conserved quantity. With the exception of the pressure gradient force, all of the terms in the inviscid Euler equations, formulated along Lagrangian surfaces, can be expressed as fluxes. The fluxes are evaluated through the Lin & Rood (1996) FV advection scheme, extended to the cubed sphere by Putman and Lin (2007), based on a two-dimensional combination of one-dimensional flux operators. The operators in the most recent version of FV3 all use the piecewise-parabolic method (Collella and Woodward 1984), although this method gives great flexibility in choosing the subgrid reconstructions used to evaluate the fluxes.

Some understanding of finite-volume methods is necessary to understand the operators. A true finite-volume method treats all variables as cell-mean values, instead of point values. (Fluxes and winds are treated as face-mean values.) To compute the fluxes, a *reconstruction* function describing an approximation to the (unknown) subgrid distribution of each variable is needed. In PPM, this reconstruction is a quadratic function, specified by the cell mean value and the values on the interfaces of the cells. The cell-interface values are derived by computing a high-order interpolation from the cell-mean values to point values on the interfaces. The reconstructions can then be altered by various means to enforce different shape preservation properties, such as monotonicity or positivity, by looking at the cell-mean and cell-interface values of the surrounding cells. This process, called *limiting* is a *non-linear* method, allowing us to create a high-order monotone solution without violating Godunov's Theorem. (This should not be confused with flux-corrected transport, or FCT, a simpler means of enforcing monotonicity)

Monotonicity means that the reconstructions are chosen to ensure that no new extrema appear in the variable, in the absence of flow convergence or divergence. This is a property of scalar advection in real flows, and prevents the appearance of unphysical oscillations, or over- and under-shoots (which can create spurious condensates or unphysically rapid chemical reactions). Monotonicity also prevents the appearance of negative values in positive-definite mass scalars, which can rapidly lead to numerical instability. However, the stronger the enforcement of monotonicity, the more diffusive the solution is, so other techniques, such as schemes which only enforce positivity, or simple reconstruction filters.

Here we briefly describe three PPM operators, all formally the same fourth-order accuracy but with different reconstruction limiters: An unlimited (also called linear) "fifth-order" operator (hord = 5), an unlimited operator with a 2dx filter (hord = 6), and the monotone Lin 2004 operator (hord = 8). We also describe the optional positive-definite limiter which can be applied to methods 5 and 6, which is essential for scalar advection (hord_tr). **The number indicated by hord is only the selection amongst a series of schemes**, similar to how different numbers in

the physics correspond to different schemes. **They do not change the order of accuracy of the advection, only the diffusivity and shape-preserving characteristics.**

Hord = 5 is the “linear” or “unlimited” scheme originally devised by Collella and Woodward, modified with a weak 2dx filter: if the cell-interface values show a 2dx signal (that is, that the interpolated interface values on opposing sides of a cell switch between greater than and less than the cell-mean value) the flux is evaluated as a first-order upwind value. The 2dx filter is necessary to suppress 2dx noise; without this filter the solver is more prone to instability. Hord = 6 uses a much stronger 2dx filter: the hord = 5 method is extended by reverting to first-order upwind flux if the difference in cell-interface values exceeds the mean of the two interface values by a tunable threshold (1.5x by default).

The optional positive-definite limiter is activated by adding a minus sign to hord (for values ≤ -6). This forces all cell-interface values to be no more than 0, ensuring a positive-definite flux. The positive-definite hord = 5 further applies the positive-definite constraint of Lin and Rood (1996).